Enabling reliable cardiovascular simulations via uncertainty quantification

Justin Tran[2], Jongmin Seo[1], Casey Fleeter[3], Daniele Schiavazzi[4], Andrew Kahn[5], Alison Marsden[1]

[1]Bioengineering and Pediatrics, Stanford University, [2]Mechanical Engineering, Stanford University, [3]ICME, Stanford University, [4]ACMS, University of Notre Dame, [5]Department of Medicine, UCSD

Introduction and Motivation

Patient-specific computational cardiovascular models are successfully employed in a wide range of clinical applications from disease diagnosis, surgical planning, and medical device design. Results, however, are often reported as deterministic, neglecting variations that could occur due to uncertain input parameters. Examples of uncertain input parameters include noisy and limited resolution medical image data, clinical measurement of patient data, or population variability in results published in the literature. Systematic quantification of uncertainties is a necessary step towards clinical adoption of computational tools.

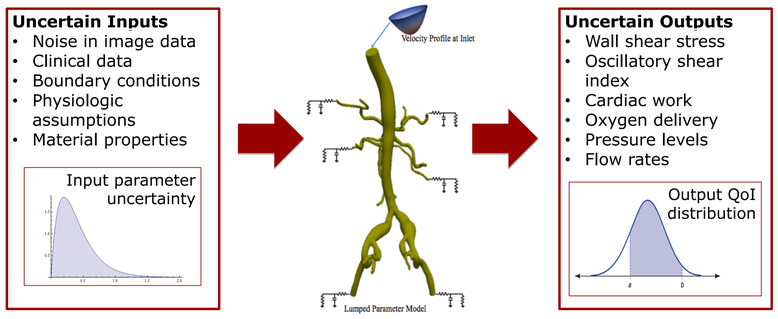

Additionally, these models often require laborious hand tuning of parameters to ensure simulation outputs mimic patient-specific behavior. This time-consuming process requires expert user knowledge, is difficult to systematically reproduce, and prevents extension of computational tools to large patient cohorts. We present a suite of efficient and automated tools for 1) assimilation of uncertain clinical data into lumped parameter boundary conditions, and 2) propagation of uncertainty to assign confidence intervals to simulations and predictions.

Patient-specific cardiovascular flow simulations

Patient-specific cardiovascular modeling in the open source software SimVascular (www.simvascular.org) consists of several steps that start from medical image data and ends with solving the incompressible Navier-Stokes equations on a finite element mesh.

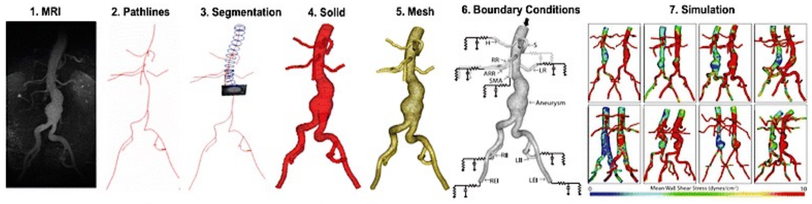

Image reproduced from Updegrove et. al. and Les et. al.

First, centerline paths are generated for all vessels of interest. Next, the vessel cross sections are segmented along these centerline paths. These segmentations are then lofted together to form a solid model, which is then typically meshed into tetrahedral finite elements. Boundary conditions are then applied before solving the incompressible Navier-Stokes. Once the simulations finish, they can be post-processed to compute hemodynamic quantities of interest. Typical quantities of interest from these simulations include pressure, flow, wall shear stress, and wall strain.

Data assimilation and parameter estimation in multi-scale simulations of coronary flow

State-of-the-art cardiovascular simulations employ lumped parameter networks (LPN) to specify boundary conditions. Manual tuning of these parameters is required to ensure simulations accurately model patient physiology, but this process is time-consuming, operator dependent, and prevents extension to large patient cohorts. We thus adopt a Bayesian perspective, treating inputs as random variables and sampling parameter sets which produce results consistent with data. We typically use a combination of patient-specific and literature data, summarized below:

Target	Description	Uncertainty	Weight	Specific/Literature
Min P_{ao}	Diastolic aortic pressure	10% measured	1	Patient-specific
Max P_{ao}	Systolic aortic pressure	10% measured	1	Patient-specific
Mean P_{ao}	Mean aortic pressure	10% measured	1	Patient-specific
$Aor - Cor$ flow split	Percentage of cardiac output to coronaries	10% mean	1	Literature
$Stroke\ Volume$	Blood volume ejected each heart contraction	10% measured	1	Patient-specific
Mean P_{pulm}	Mean pulmonary pressure	3.3 mmHg	2	Literature
$Ejection\ Fraction$	Percentage of LV blood volume ejected per contraction	10% measured	1	Patient-specific
Mitral E/A ratio	Ratio of early to late flows into the LV	20% measured	2	Patient-specific
Mitral valve open %	Percentage of cardiac cycle that mitral valve is open	15% measured	2	Patient-specific
Aortic valve open %	Percentage of cardiac cycle that aortic valve is open	15% measured	2	Patient-specific
Pulm valve open %	Percentage of cardiac cycle that pulmonary valve is open	15% measured	2	Patient-specific
Max $P_{rv} - P_{ra}$	Systolic pressure difference between the RV and RA	25% measured	2	Patient-specific
Mean P_{ra}	Mean RA pressure	40% measured	1	Patient-specific
L_{cor} peak ratio	Left coronary peak flow ratio in diastole vs. systole	0.8	1	Literature
L_{cor} total ratio	Left coronary flow volume ratio in diastole vs. systole	2.53	1	Literature
L_{cor} 1/3 FF	Percentage of left coronary flow volume in first 1/3 of cardiac cycle	0.02	1	Literature
L_{cor} 1/2 FF	Percentage of left coronary flow volume in first 1/2 of cardiac cycle	0.03	1	Literature
R_{cor} peak ratio	Right coronary peak flow ratio in diastole vs. systole	0.3	1	Literature
R_{cor} total ratio	Right coronary flow volume ratio in diastole vs. systole	1.08	1	Literature
R_{cor} 1/3 FF	Percentage of right coronary flow volume in first 1/3 of cardiac cycle	0.07	1	Literature
R_{cor} 1/2 FF	Percentage of right coronary flow volume in first 1/2 of cardiac cycle	0.07	1	Literature

Assuming a Gaussian likelihood function, we sample parameter sets from the posterior distribution using and adaptive Markov Chain Monte Carlo (MCMC). We can then use Nelder-Mead optimization of the maximum posterior parameter estimate to solve for the optimally tuned parameters which match patient data.

Our workflow is summarized below:

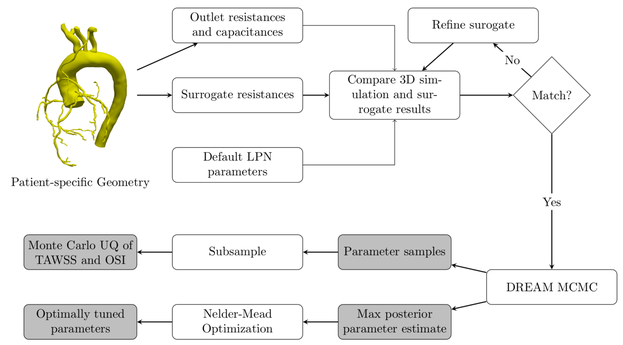

We used this framework to tune the input parameters in seven patients exhibiting a wide range of different anatomies and physiologic targets, and matched the data within their specified uncertainty, summarized below for the maximum aortic pressure, stroke volume, and ejection fraction:

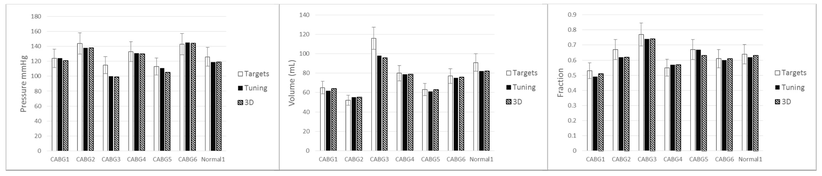

The average percent difference between simulated results and patient targets was 7.6%, under the average of 14.6% measurement uncertainty. This framework can also be easily extended to other patient anatomies and different available data. The samples produced from the MCMC are also key for forward propagation of uncertainties to model outputs.

Uncertainty quantification in simulations of coronary bypass grafts

This study aims at quantifying uncertainty in two computational results (viscous wall shear stress and mechanical wall strain) as a result of uncertainties in the LPN boundary conditions and wall material properties. We also developed a *stochastic submodelling* approach to simulate only our region of interest (bypass grafts) to alleviate the computational burden of running the many required full multiscale simulations for uncertainty quantification. This submodelling relies on re-parameterizing the velocities and pressures of the graft submodel in terms of the full model parameters, and using sparse regression to compute their relationship. This reduced the cost of running simulations by over an order of magnitude.

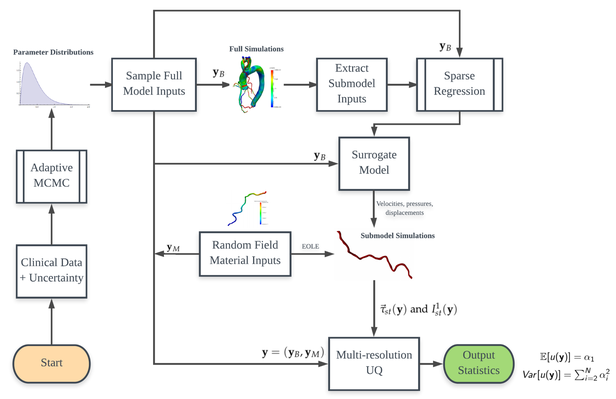

We then used a multi-resolution approach to uncertainty propagation, which extends the generalized polynomial chaos expansion. We used this framework to compute output uncertainties for three different patients, representing a wide range of physiological data and graft geometries.

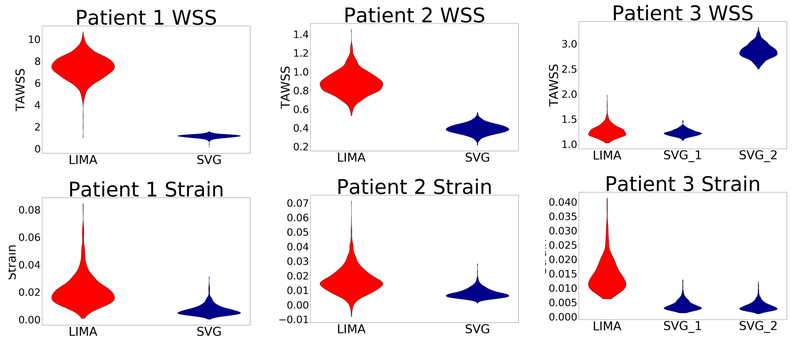

Analyzing the probability distributions for the wall shear stress (WSS) and wall strain, we see that WSS is relatively well estimated in the presence of input uncertainties as there are clear differences in the distributions between LIMA and SVG grafts. Wall strain, on the other hand, is poorly estimated with the distributions bleeding into one another. This information is key for determining which simulated outputs are most reliable to use in the clinic for affecting patient care.

Multi-fidelity framework for uncertainty quantification in cardiovascular simulations

Monte Carlo approaches will reliably converge to the true value for any quantity of interest, but the large number of simulations needed for this convergence is untenable for full model simulations. Multi-level and multi-fidelity approaches aim to reduce variance compared to that obtained when using the same number of simulations with Monte Carlo.

A pilot study using three available model fidelities compared the results of six multilevel (ML), multifidelity (MF) and multilevel-multifidelity (MLMF) methods, summarized below, to standard Monte Carlo approaches. The full UQ study currently in progress utilizes the MLMF framework, below far right:

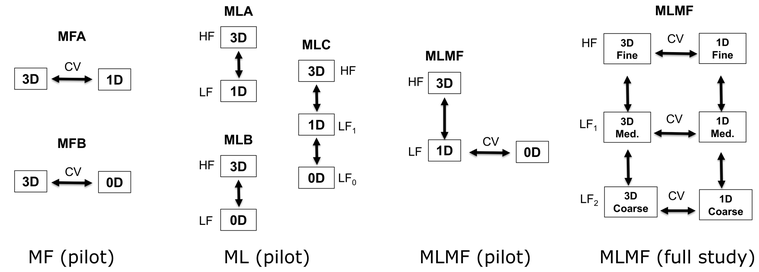

Three fidelity levels of the same healthy coronary model geometry are shown below. The Hughes and Lubliner 1D formulation with a linear constitutive equation is used in our 1D solver, while the 0D model is a full-model LPN.

Uncertainty quantification was performed using steady inlet flow, with ten resistance boundary conditions, sampled from uniform distributions about means tuned to physiologic waveforms, as the uncertain parameters. Global (steady state flow and pressure values at outlets) and local (various WSS quantities) served as the quantities of interest for the exploratory pilot study.

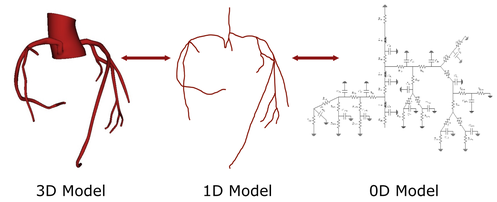

Solver	No. Simulations
3D	100
1D	2000
0D	10 000

Right: Simulations of each fidelity for pilot study.
Below: Cost of each varying fidelity solver

Solver	Cost (1 simulation)	Effective Cost (No. 3D Simulations)
3D	317 CPU hrs	1
1D	5.65 min	3E-4
0D	19 sec	1.7E-5

The cost of each method, using all simulation results of the pilot study, are compared in the table below left. Extrapolation was used to determine the number of simulations of each fidelity level needed to obtain 1% accuracy for flow at a circumflex artery subbranch outlet, results shown below right.

Method	Effective Cost (3D Simulations)	No. 3D Simulations	No. 1D Simulations	No. 0D Simulations
MC	100	100	–	–
MFA	100.6410	100	2 000	–
MFB	100.1644	100	–	10 000
MLA	100.6410	100	2 000	–
MLB	100.1644	100	–	10 000
MLC	100.8037	100	2 000	9 900
MLMF	100.8037	100	2 000	9 900

Method	Effective Cost (3D Simulations)	No. 3D Simulations	No. 1D Simulations	No. 0D Simulations
MC	55 465	55 465	–	–
MFA	30	7	71 640	–
MFB	337	319	–	1 061 964
MLA	92	47	140 351	–
MLB	464	442	–	1 293 914
MLC	42	32	9 862	395 160
MLMF	40	31	6 886	405 863

The accuracy of each quantity of interest for each method is shown in the graph below left. Accuracy is defined as $(6\sqrt{(\mathrm{Var}[Q]}\ ))/\mathrm{E}[Q]$ for each quantity of interest Q (ratio of confidence interval length to expected value).

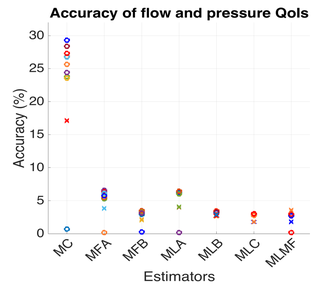

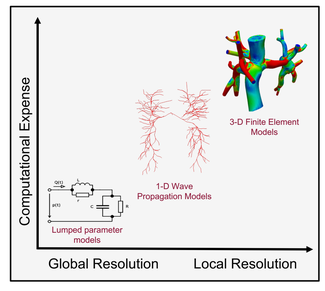

The promising results of the pilot study have led to a current study comparing the MLMF methods on healthy and diseased models, utilizing more realistic boundary conditions as well as a wider range of both quantities of interest and uncertain parameters. The advantages of this method for our application are shown in the graph above right.

References and Acknowledgments

Funding from NIH-NIBIB R01 EB018302, RO1 RHL123689A.
Computational resources from XSEDE and Stanford Research Computing Center's Sherlock cluster.

Updegrove A, Wilson NM, Merkow J, et. al.. SimVascular - An open source pipeline for cardiovascular simulation. *Ann Biomed Eng*. 2016; doi:10.1007/s10439-016-1762-8.
Tran, Justin S., et al. "Automated tuning for parameter identification and uncertainty quantification in multi-scale coronary simulations." Computers & fluids 142 (2017): 128-138.
Les A, Shadden S, Figueroa C, et. al.. Quantification of Hemodynamics in Abdominal Aortic Aneurysms During Rest and Exercise Using Magnetic Resonance Imaging and Computational Fluid Dynamics. *Ann Biomed Eng*. 2010; doi: 10.1007/s10439-010-9949-x
Geraci G, Eldred M, Iaccarino G. A multifidelity multilevel Monte Carlo method for uncertainty propagation in aerospace applications. *19th AIAA Non-Deterministic Approaches Conference*. 2017.
T.J.R. Hughes and J. Lubliner. On the One-Dimensional Theory of Blood Flow in the Larger Vessels. *Mathematical Biosciences*. 18(1-2) (1973), 161-170.